

Avoiding Medication Conflicts for Patients with Multimorbidities*

Andrii Kovalov¹ and Juliana K. F. Bowles²

¹ German Aerospace Center (DLR), Simulation and Software Technology
`andrii.kovalov@dlr.de`

² School of Computer Science, University of St Andrews
KY16 9SX, St Andrews, UK
`jkfb@st-andrews.ac.uk`

Abstract. Clinical pathways are care plans which detail essential steps in the care of patients with a specific clinical problem, usually a chronic disease. A pathway includes recommendations of medications prescribed at different stages of the care plan. For patients with three or more chronic diseases (known as *multimorbidities*) the multiple pathways have to be applied together. One common problem for such patients is the adverse interaction between medications given for different diseases. This paper proposes a solution for avoiding medication conflicts for patients with multimorbidities through the use of formal methods. We introduce the notion of a pharmaceutical graph to capture the medications associated to different stages of a pathway. We then explore the use of an optimising SMT solver (Z3) to quickly find the set of medications with the minimal number and severity of conflicts which is assumed to be the safest. We evaluate the approach on a well known case of an elderly patient with five multimorbidities.

1 Introduction

There is an increasing number of people in Europe and in the UK with three or more long term conditions, also known as *multimorbidities*. In Europe, the current number of people with multimorbidities is estimated at around 50 million [7], and in the UK it is currently around 1.9 million [8]. Chronic diseases often develop simultaneously in response to common risk factors such as smoking, diet, ageing, and inactivity [3, 18]. The four most common chronic diseases are cancer, chronic obstructive pulmonary disease (COPD), coronary heart disease, and diabetes.

Despite the growing prevalence of chronic disease, strategies for improving the management of patients with co-morbidities remain under-explored. In clinical settings processes are complex and necessarily rely on a range of interacting social agents including physicians, administrators and patients who in turn are influenced by a number of social, technical and organisational factors. This complexity can result in variation in how physicians practice, appropriate care is

* This work is partially supported by EPSRC grant EP/M014290/1.

documented, and healthcare costs are managed [19]. To reduce inconsistencies clinical guidelines have emerged which are based on the best existing evidence and which aim to support clinical staff and improve the quality of healthcare. There are currently around 180 clinical guidelines published by the National Institute for Clinical Excellence (NICE)³. However, current guidelines almost entirely focus on single conditions. As a result, applying multiple guidelines to a patient may potentially result in conflicting recommendations for care. In the UK and Europe there have been calls for improved integration of existing guidelines for patients with multimorbidities [12, 7], but this is still very much an open problem.

To encourage the translation of national guidelines into local protocols, and subsequently clinical practice, *clinical pathways* have been developed. Clinical pathways are care plans which detail essential steps in the care of patients with a specific clinical problem, usually a chronic disease [6]. These pathways frequently use graphical descriptions of evidence and options, and are typically represented in a single or a series of flow charts [9, 16]. A pathway includes recommendations of medications prescribed at different stages of the care plan. One common problem for patients with multiple long term conditions is the adverse interaction between medications given for different diseases. We concentrate on such medication interactions in this paper, and introduce a notion of a *pharmaceutical graph* for a chronic disease as a formal representation of the medications and medication groups underlying a clinical pathway. These pharmaceutical graphs and the information on known adverse medication conflicts are used to answer the fundamental question of *what is the best care plan and what are the most effective medications in the treatment of patients with multimorbidities?*

The main contribution of our paper is a solution for avoiding medication conflicts for patients with multimorbidities through the use of formal methods. We explore the use of an optimising SMT solver (Z3) to quickly find the set of medications with the minimal number and severity of conflicts which is assumed to be the safest. We evaluate our approach with a very well known case in the medical domain taken from [5] of a hypothetical 79-year-old woman with five long term conditions: hypertension, diabetes mellitus (type 2), osteoarthritis, osteoporosis and chronic obstructive pulmonary disease (COPD).

This paper is structured as follows. We start by setting the context of our work and describing some of the most relevant existing related work in Section 2. Section 3 introduces our notion of *pharmaceutical graph* which is extracted for each disease from the corresponding documented clinical pathways. Since our approach makes use of SMT solvers, in Section 4 we describe how to derive logical formulae from the pharmaceutical graphs and medication conflicts. Our approach is evaluated in Section 5, and its performance is compared with other possible solutions. We conclude our paper in Section 6 with a discussion of current limitations and ideas for future work.

³ NICE www.nice.org.uk

2 Related Work

The National Institute for Clinical Excellence (NICE) has a considerable number of documented pathways for care and treatment of diseases including the most common chronic conditions such as cancer, chronic obstructive pulmonary disease, coronary heart disease, diabetes, hypertension, osteoarthritis, chronic heart failure, chronic kidney disease, depression and so on. The pathways are nonetheless essentially single-disease descriptions, and there is little advice on how to combine or consider two or more of them when addressing the needs of patients with multimorbidities. Better approaches for integrated care are lacking and urgently needed.

In clinical practice, the main suggestion is to enrich pathways with additional information. This includes recommendations regarding certain comorbidities, discussion of benefits and risks of treatment, and advice on treating elderly patients [5, 10, 15].

More recently, the possibility of developing automated techniques for integrated care has drawn attention from the computer science community. Several approaches were introduced with the aim of formalising and merging existing single-disease pathways to produce a treatment advice for patients with multiple diseases. A few different approaches are described next.

One approach is to model pathways using ontologies [2, 1, 11]. The authors focus on eliminating duplicated tasks across different pathways (e.g. blood tests), reusing the results of common activities and avoiding medication conflicts. Although ontologies are an expressive modelling instrument, the automatic merging of multiple ontologies seems problematic.

Another approach also making use of ontologies is carried out under the project GLARE [17] which instead of automatically merging guidelines provides an interactive interface for clinicians to analyse multiple guidelines at different levels of detail.

A different approach is shown in [13] where the authors use rewriting rules to eliminate conflicts in merged guidelines.

A research somewhat similar to ours is presented in [20]. The authors encode individual pathways as sets of formulas in first-order logic and use an automated theorem prover to find a combined treatment plan. The medication conflicts are also represented as logical expressions, and if a conflict is detected then special ‘Revision Operators’ are invoked that rewrite fragments of logical expressions so that the conflict may disappear. These operators correspond to some medical actions that are performed to resolve a conflict (such as the co-prescription of an additional medication). If the Revision Operators cannot resolve the conflict then the algorithm fails to produce any treatment suggestion.

Our approach is similar in that we also transform medical knowledge into logical expressions and use a tool (an SMT solver rather than a theorem prover) to automatically produce treatment advice. However, unlike [20], we take into account the medical utility of the produced advice, and try to generate not just any possible advice, but the *best and safest* possible recommendation with respect to given knowledge on medication conflicts. The advantage of our approach

is that we will be able to provide a solution even if certain conflicts cannot be avoided. In the case of patients with 5 or more multimorbidities, it is likely that some conflicts are indeed unavoidable and the aim in such cases is to reduce it to a more favourable alternative.

3 Pharmaceutical Graphs

The goal of our research is to help clinicians make more informed decisions, and more concretely medication prescriptions that result in the minimal amount and severity of medication conflicts. In other words, we would like to find the safest set of medications for a patient with multiple diseases. We take into account three types of conflicts:

- Drug-drug conflicts (when two medicines taken together have some negative effect or decrease effectiveness of each other),
- Drug-disease conflicts (e.g. any medication that increases blood pressure conflicts with hypertension),
- Drug-patient conflicts (personal medication intolerances and allergies).

To detect and avoid these conflicts we need to know which medications can be prescribed according to the clinical pathway. As the pathways are written informally in natural language, we first need to create a formal representation of a pathway. We developed a representation which we call a *pharmaceutical graph* which is a directed acyclic graph with one initial node where the nodes represent medication prescriptions. Normally a clinical pathway suggests several alternative medications for the doctor to choose from. Likewise, in our pharmaceutical graph a node can contain multiple medications one of which should be prescribed.

Examples of pharmaceutical graphs for Diabetes Melitus and Hypertension are shown in Fig. 1. For clarity we included a dummy initial node with the disease name that does not contain any medications. For Diabetes this dummy node is also necessary for branching.

Often the pathways recommend not a single medication, but a medication group such as *Sulfonylurea*. In this case we include all individual Sulfonylureas into the node. In Fig. 1 for the medication groups the number of individual medications is shown in brackets. It is important to list all the medications in a group because they might have different conflicts.

Pharmaceutical graphs capture two features of the original pathways: the structure (ordering and branching) and the advised medications. A maximal path in a pharmaceutical graph (from the initial node to a leaf node) represents a choice of medications in a treatment plan. For example, one maximal path (prescription) in the Diabetes graph could be Metformin, Sitagliptin, Insulin.

Having such maximal paths for multiple pharmaceutical graphs, we can detect the above mentioned three types of drug-related conflicts in them.

Additionally, we consider the severity of medication conflicts and how suitable given medications are known to be for the treatment of certain conditions.

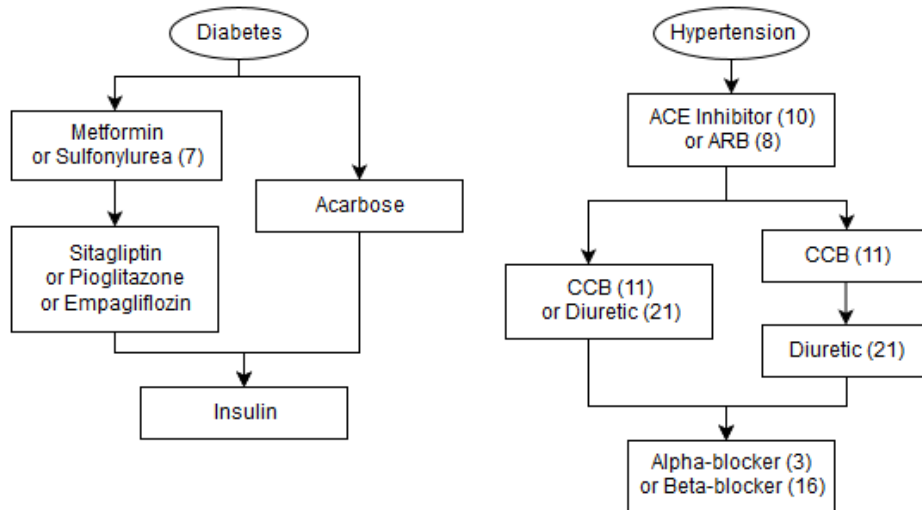


Fig. 1. Pharmaceutical graphs for Diabetes and Hypertension based on NICE pathways. ACE - angiotensin-converting-enzyme, ARB - angiotensin receptor blocker, CCB - calcium channel blocker. The numbers in brackets represent the number of individual medications in a group.

To distinguish between major and minor conflicts we assign negative integer scores to them. Similarly, we assign positive scores to the medications depending on their efficacy. Our aim is to obtain a treatment plan with the highest score considered to be the safest and most effective for a particular patient. In this way we can produce a reasonable treatment advice even when certain conflicts cannot be avoided. Note that it is outside the scope of our work to assign scores to drugs. The safety and effectiveness of drugs is subject to evaluations carried out in clinical pharmacology and biopharmaceutics data research. Pharmaceutical companies assign scores to drugs based on evidence-based safety and effectiveness.

Our question is, therefore, formulated as follows: *Having a set of pharmaceutical graphs and a set of conflicts, what are the prescriptions (maximal paths) for all the graphs that result in the highest total score?* The total score here is the sum of the positive scores of all the chosen medications and the negative scores of the triggered conflicts. Thus, we have transformed a medical problem into a computation problem. In the next section we show how we tackle it using an optimising SMT (satisfiability modulo theories) solver.

4 Transformation into an SMT problem

SMT solvers extend the boolean satisfiability problem (SAT) with additional theories for integers, arrays, etc. Given a predicate over boolean and integer variables, an SMT solver can check if there exists an assignment for the variables

that turns the predicate into true. If such assignments exist, an SMT solver produces one of these assignments.

We are using an SMT solver Z3 developed at Microsoft research [14] which has a built-in optimisation functionality. In other words, Z3 can produce not just any satisfying assignment but an optimal assignment with respect to some objective function [4].

The core idea of our work is to translate the pharmaceutical graphs and the medication conflicts into logical and arithmetic expressions and use Z3 to generate maximal paths for the pharmaceutical graphs with the highest possible total score.

There are two parts of this translation that we would like to discuss in the following subsections - the representation of the graph structure and the representation of the medications and conflicts along with their scores.

4.1 Translating the graph structure

We need to translate the structure of the pharmaceutical graphs into logical expressions and make Z3 produce a maximal path.

First, we shrink the original pharmaceutical graph in order to reduce the number of variables we need to represent its structure (see Fig. 2).




Fig. 2. Original pharmaceutical graph with branches shown in grey (left) and the shrunk branch graph (right).

To do this, we divide the nodes of the pharmaceutical graph into different *branches*. A branch here is the longest possible sequence of nodes where there is a path between the first and the last node, and along this path all the nodes except the last one have exactly one outgoing edge, and all the nodes except the first one have exactly one incoming edge. The branches are graphically shown in Fig. 2 with grey regions.

An important property of such a branch is that for any maximal path in the pharmaceutical graph either the whole branch is on the path or the whole branch is not on the path. Therefore, we can treat a maximal path as a sequence of

branches. If we substitute each branch with a single node, we will get a shrunk branch graph (see Fig. 2, right) that will have the same maximal paths as the original pharmaceutical graph (in terms of branch sequences).

Our idea is to create a boolean variable for each branch and use Z3 to assign these branch variables to true or false so that there will be one maximal path consisting of ‘true’ branches and all the rest will be ‘false’. To do this, we specify logical assertions that would filter out all the invalid assignments of the branch variables (that do not make one maximal path). We tried and compared two different approaches to produce these logical assertions.

The first and a more naive approach, which is taken in [20], is to explicitly enumerate all the maximal paths. Illustrating this for the graph in Fig. 2, we obtain the following logical statement:

$$(Br0 \wedge Br1 \wedge Br3 \wedge Br5 \wedge \neg Br2 \wedge \neg Br4 \wedge \neg Br6) \vee \\ (Br0 \wedge Br1 \wedge Br4 \wedge Br5 \wedge \neg Br2 \wedge \neg Br3 \wedge \neg Br6) \vee \dots$$

For the second approach, we designed a set of *branching rules* that are applied to every node in the graph and produce some logical statements. The rules are as follows:

1. The initial branch is always true
2. If the current branch is true then one of its direct successors is true and the other direct successors are false (if there are direct successors)
3. If all the direct predecessors of the current branch are false then the current branch is false (if there are direct predecessors)

For example, for branches $Br0$, $Br1$ and $Br4$ in Fig. 2 the following statements will be generated.

$$\begin{array}{ll} Br0 & \text{(rule 1 for Br0)} \\ Br0 \implies (Br1 \wedge \neg Br2) \vee (\neg Br1 \wedge Br2) & \text{(rule 2 for Br0)} \\ \\ Br1 \implies (Br3 \wedge \neg Br4) \vee (\neg Br3 \wedge Br4) & \text{(rule 2 for Br1)} \\ \neg Br0 \implies \neg Br1 & \text{(rule 3 for Br1)} \\ \\ Br4 \implies (Br5 \wedge \neg Br6) \vee (\neg Br5 \wedge Br6) & \text{(rule 2 for Br4)} \\ \neg Br1 \wedge \neg Br2 \implies \neg Br4 & \text{(rule 3 for Br4)} \end{array}$$

We are interested in the variable assignments that turn these statements true for all the nodes of a branch graph, so we pass these statements as assertions to Z3, and get an assignment representing a maximal path.

Proposition 1. *The three branching rules ensure that exactly one maximal path in a branch graph will consist of true nodes, and all the other nodes will be false.*

Proof. First we will show that there will be a maximal path consisting of true nodes. The initial node in the branch graph will be true (rule 1). This means one of its direct successors will also be true (rule 2), one of the direct successors of this true successor will be true as well (rule 2) and so on until a leaf node for which rule 2 will not produce any statements. Therefore, there will be a path starting in the initial node and finishing in a leaf node (i.e. a maximal path) consisting of true branches.

We will now prove by contradiction that no other node can be true. Let us assume there is a node A which is true and does not belong to this maximal path. A is not the initial node (because otherwise it would belong to a maximal path by definition). Therefore, A has some direct predecessor B . This predecessor may or may not belong to the maximal path. If it belongs to the maximal path it should have another direct successor C apart from A which is true and lies on the path. Consequently, A cannot be true because according to rule 2 all the other direct successors of B apart from C should be false.

If no direct predecessor of A belongs to the maximal path, at least one of them should be true (otherwise A would be false according to rule 3). This true predecessor should also have a true predecessor (rule 3) and so on, so there will be a chain of true predecessors. Our graph is acyclic and has only one initial node, so if we follow the edges backwards from any node we will always come to the initial node which lies on the maximal path. Therefore, in this chain of true predecessors there will be some node D which has a direct predecessor that belongs to the chosen maximal path. As we already showed above, this node D cannot be true. We have come to a contradiction, so a true node A that does not belong to the maximal path cannot exist.

□

It is hard to tell whether the naive approach or the approach with the branching rules is more suitable for our task because it depends on how Z3 processes the statements internally. We compare both approaches in Section 5.

4.2 Translating medications, conflicts and scores

Aside from having a maximal path in the pharmaceutical graph we also need to choose the individual medications in every node along this path. Therefore, if a node belongs to a ‘true’ branch we want to pick one medication. Additionally, as we already mentioned, every medication has an integer score which we also need to pass to Z3.

For example, in the Diabetes graph shown in Fig. 1 there is a node ‘Metformin or Sulfonylurea’. Let us assume this node has the identifier $N1$ and belongs to a branch $Br1$, the score of Metformin is 100 and the score of Sulfonylurea is 70. Sulfonylureas are a medication group rather than a single medication, but we ignore this for the moment.

The assertions we generate for this node are shown below.

$$Br1 \implies (Metformin \wedge \neg Sulfonylurea) \vee (\neg Metformin \wedge Sulfonylurea)$$
$$\neg Br1 \implies \neg Metformin \wedge \neg Sulfonylurea$$
$$N1_Score = \text{if } (Metformin) \text{ then } 100 \text{ else if } (Sulfonylurea) \text{ then } 70 \text{ else } 0$$

For every medication we create a boolean variable that denotes whether it is prescribed or not. If the node lies on the chosen maximal path ($Br1$ is true) then one medication out of two will be set to true, otherwise both will be false.

We also introduce an integer variable $N1_Score$ which will get different value depending on which medication is picked (Z3 supports integers and if-then-else expressions). We create score variables for every node of the pharmaceutical graph and then sum them to get the total medication score.

As for the conflicts, we represent them as integer variables in the following way.

$$Conflict1_Score = \text{if } (Nadolol \wedge Diabetes) \text{ then } -5000 \text{ else } 0$$
$$Conflict2_Score = \text{if } (Sitagliptin) \text{ then } -1000 \text{ else } 0$$
$$Conflict3_Score = \text{if } (Metformin \wedge Acarbose) \text{ then } -100 \text{ else } 0$$

Here $Conflict1$ is a major drug-disease conflict, $Conflict2$ is a moderate allergy and $Conflict3$ is a minor drug-drug conflict.

The total score is then the sum of the medication scores (which are non-negative) and the conflict scores which are negative when a conflict is triggered.

To put it all together, we start with a set of pharmaceutical graphs corresponding to the diseases of a particular patient and a set of conflicts which may arise. Then we generate Z3 code as explained above and run Z3 with the total score as the objective function. The output of Z3 are the assignments of the branch and medication variables that correspond to a treatment plan with the highest score.

In the next section, we show an example of using this system and discuss the results.

5 Evaluation

To evaluate our approach, we model a medically well known case from [5] of a hypothetical 79-year-old woman with five diseases: hypertension, diabetes mellitus (type 2), osteoarthritis, osteoporosis, and chronic obstructive pulmonary disease (COPD).

5.1 Medical data collection

We used two sources of medical information. We use the clinical pathways for the diseases as documented by NICE and an online portal Drugs.com⁴ to get the data on the medication groups and conflicts.

⁴ <http://www.drugs.com>

NICE pathways are available in the form of flowcharts with supporting text. A fragment of the diabetes pathway is shown in Fig. 3. On the left side it shows the graphical representation of the current aspect of the pathway (in the example the therapy for lowering blood-glucose for patients with type 2 diabetes). Details for the current node being highlighted (Metformin) are shown in the text on the right side.

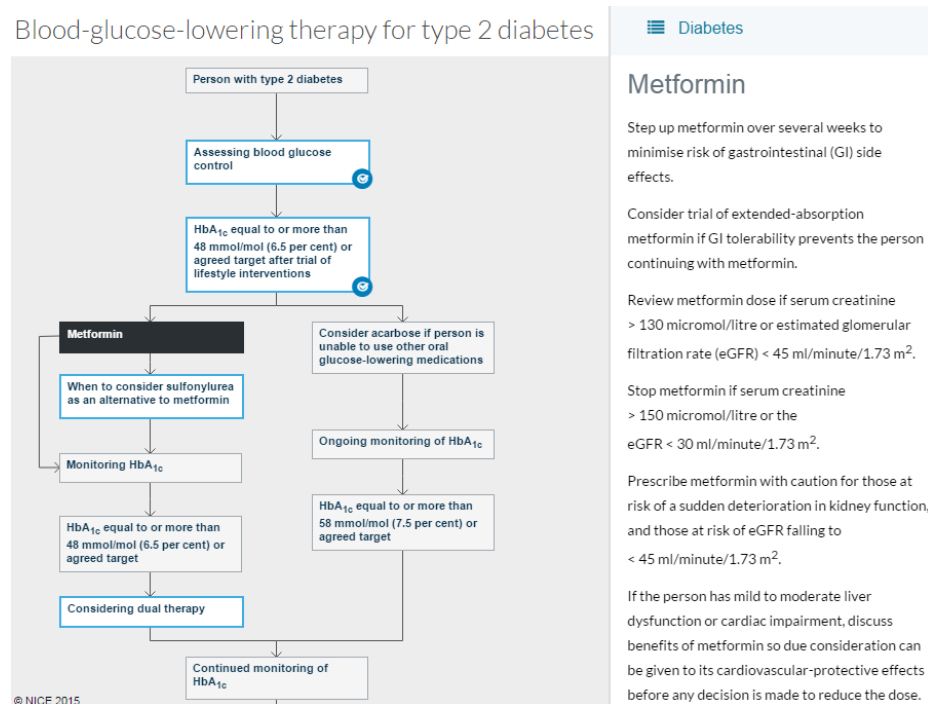


Fig. 3. Fragment of a NICE clinical pathway for type 2 diabetes.

Given such a documented pathway for a long term condition, the first step is to extract the information about the medications and transform it into our pharmaceutical graph. We did this manually for all five long term conditions considered. The resulting pharmaceutical graphs are shown in Fig. 1 and Fig. 4.

When pathways recommended a medication group, we expanded it based on the data from Drugs.com. In total for all five diseases, our pharmaceutical graphs contain 127 medications.

To get all the possible drug-drug conflicts, we used an interaction checker of Drugs.com⁵. To get the drug-disease conflicts, we examined the medication profiles on Drugs.com. We only considered drug-drug conflicts where both conflicting medications appear in our pharmaceutical graphs and drug-disease conflicts

⁵ http://www.drugs.com/drug_interactions.html

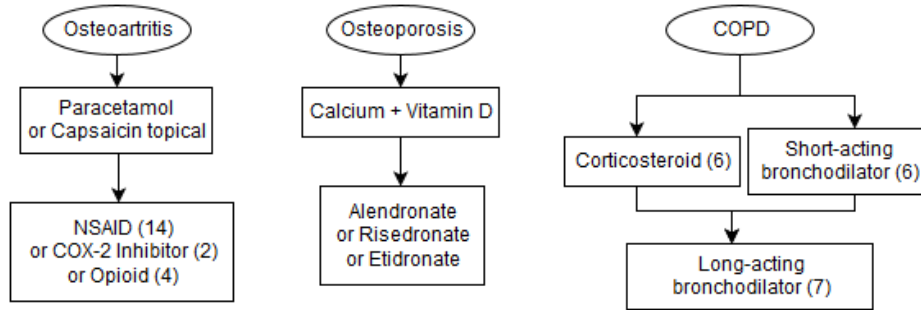


Fig. 4. Pharmaceutical graphs for Osteoarthritis, Osteoporosis and Chronic Obstructive Pulmonary Disease (COPD) based on NICE pathways. NSAID - nonsteroidal anti-inflammatory drug, COX-2 - Cyclooxygenase-2.

where the medication appears in the pharmaceutical graphs and the disease is one of the considered five diseases. The summary of the conflicts is shown in Table 5.1. As we can see, the number of conflicts is far too large to analyse manually. We did not include any drug-patient conflict in our experiment.

Table 1. Conflicts summary

	Major	Moderate	Minor	Total
Drug-drug	270	3038	178	3486
Drug-disease	47	33	0	80
Total	317	3071	178	3566

As for the score values, we assigned medication scores according to the rating on Drugs.com in the range from 0 to 100. For minor, moderate, and major conflicts we assigned values -100, -1000, -5000 respectively.

Ideally we would like medical experts to assign score values, however, for our experiment we just assigned values that seemed reasonable to us and in accordance to the pharmaceutical ratings mentioned above.

5.2 Experiment results

To run this experiment, we created a simple application which takes the pharmaceutical graphs and the conflicts as an input, generates the corresponding Z3 code, runs it, parses Z3 output and renders the resulting optimal set of maximal paths using Graphviz.

The output of our application for this experiment (with 5 diseases, 127 medications, and 3566 conflicts) is shown in Fig. 5. As we can see, the best recommendation still causes four moderate conflicts, and the total score is dominated by the score of these conflicts. In this case the doctor might decide to exclude

some medications from the treatment, because the negative effect caused by the conflicts may outweigh the benefits.

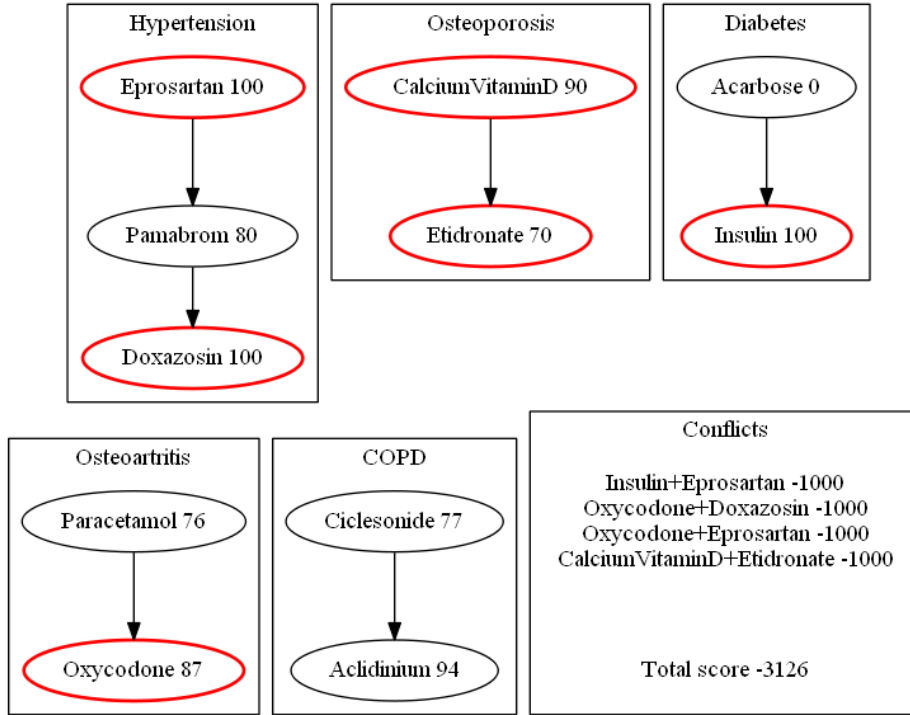


Fig. 5. Optimal medication prescriptions for a hypothetical patient with five diseases generated by Z3.

To evaluate the performance of the system, we ran the experiment on an off the shelf laptop with 16 GB RAM and a quad-core 2.6 GHz processor. We tried both approaches to produce logical assertions discussed in Subsection 4.1. The average run time with the explicit enumeration of maximal paths is 7.6 seconds, while with the branching rules it is 8.1 seconds, so the explicit enumeration works marginally faster for this case.

In our opinion, the system runs fast enough to be used in a clinical environment as a decision support tool.

The real run time is likely to be less because doctors might exclude medications that are not certified or that are known to have many negative interactions. Additionally, doctors might pick a particular medication or branch manually and therefore reduce the amount of computations required. Moreover, it is only a minority of patients that will have five or more concomitant conditions.

5.3 Other approaches

In addition to using optimising Z3, we evaluated several other approaches that turned out to be less efficient.

Before the optimising version of Z3 was released, we tried iterative SMT solving using a hill-climbing and a dichotomy approaches. In a hill-climbing approach we first obtained some initial solution with score S_0 , and then added an assertion $FinalScore > S_0$. The solver would return another solution with score $S_1, S_1 > S_0$. We then continued to get better and better solutions until we hit ‘unsatisfiable’.

In the dichotomy approach we used the fact that the best score lies in a bounded interval. It cannot be greater than the sum of all the medication scores and cannot be less than the sum of all the negative conflict scores. We calculated this interval and iteratively halved it by examining the mid-point. We used an SMT solver to check if a solution existed with the score greater or equal to this mid-point, and depending on the result took the upper or the lower half of the interval until we reached an interval containing only the highest score.

These approaches are much slower than using the built-in optimising Z3 functionality (run time is over an hour for the experiment with five diseases compared to under 10 seconds for the optimising Z3).

We also tried to find the optimal solution using A*. This worked comparatively well on smaller problems, but for the experiment with five diseases our A* implementation ran out of memory (1 Gb) after a minute of calculation. This was due to the growth of the set of visited states. We defined a state as a set of incomplete or complete paths for the input pharmaceutical graphs. Every possible combination of medications resulted in a new state, leading to a very large search space.

We conclude from this evaluation that the approach based on SMT solving is more appropriate than A* for heavily constrained problems (in our case the constraints are the numerous medication conflicts). For such problems it is hard to define a good cost under-approximation function for A* which makes the algorithm explore wrong directions.

6 Conclusions

Our work shows that a treatment advice for a patient with multimorbidities can be produced automatically using an SMT solver. This approach can be used in a clinical decision support tool as it works reasonably fast. Our work acts as a proof of concept for a tool which is being developed for clinicians.

Additionally, the approach allows us to choose a certain medication or a pathway branch manually. We believe that this kind of interactive analysis can be useful to get some insight into the different treatment options and their consequences.

Our current way of modelling the clinical pathways and underlying medication conflicts also has some limitations.

First of all, the dosages and the timing information for conflicts should be taken into account. Some interactions may arise only when the medications are taken simultaneously or if the dose is large enough.

Secondly, it is not clear how to assign numerical scores to medications in general, and especially in the case of branching. The branches may contain different number of medications, and a longer branch will likely get a higher score. Therefore, the scores should be somehow balanced.

Finally, our current approach aims at avoiding, but not resolving conflicts. For instance, sometimes it is necessary to co-prescribe additional medications in order to mitigate a certain adverse interaction, and we currently do not consider this issue.

In future work we aim to address these limitations, and eventually create a tool that will help clinicians deliver better care for an increasing number of patients suffering from multimorbidities.

References

1. Abidi, S.R.: Ontology-Based Knowledge Modeling to Provide Decision Support for Comorbid Diseases. In: Riaño, D., ten Teije, A., Miksch, S., Peleg, M. (eds.) Knowledge Representation for Health-Care, Lecture Notes in Computer Science, vol. 6512, pp. 27–39. Springer (2011)
2. Abidi, S.R., Abidi, S.S.R.: Towards the Merging of Multiple Clinical Protocols and Guidelines via Ontology-Driven Modeling. In: Combi, C., Shahr, Y., Abu-Hanna, A. (eds.) Artificial Intelligence in Medicine, Lecture Notes in Computer Science, vol. 5651, pp. 81–85. Springer (2009)
3. Barnett, K., Mercer, S., Norbury, M., Watt, G., Wyke, S., Guthrie, B.: Epidemiology of multimorbidity and implications for health care, research, and medical education: a cross-sectional study. *Lancet* 380(9836), 37–43 (2012)
4. Bjørner, N., Phan, A.D., Fleckenstein, L.: νZ - An Optimizing SMT Solver. In: Baier, C., Tinelli, C. (eds.) Tools and Algorithms for the Construction and Analysis of Systems, Lecture Notes in Computer Science, vol. 9035, pp. 194–199. Springer (2015)
5. Boyd, C.M., Darer, J., Boult, C., Fried, L.P., Boult, L., Wu, A.W.: Clinical Practice Guidelines and Quality of Care for Older Patients With Multiple Comorbid Diseases: Implications for Pay for Performance. *JAMA* 294(6), 716–724 (2005)
6. Campbell, H., Hotchkiss, R., Bradshaw, N., Porteous, M.: Integrated care pathways. *BMJ* 316(7125), 133–137 (1 1998)
7. Directorate-General for Health and Food Safety, European Commission: Conference "Which priorities for a European policy on multimorbidity?", Final Report, Brussels, 27 October (2015)
8. Long term conditions compendium of information, 3rd ed. (2012), department of Health
9. Hawley, S., Zikmund-Fisher, B., Ubel, P., Jancovic, A., Lucas, T., Fagerlin, A.: The impact of the format of graphical presentation on health-related knowledge and treatment choices. *Patient Education and Counseling* 73(3), 448 – 455 (2008)
10. Hughes, L.D., McMurdo, M.E.T., Guthrie, B.: Guidelines for people not for diseases: the challenges of applying UK clinical guidelines to people with multimorbidity. *Age and Ageing* 42(1), 62–69 (2013)

11. Jafarpour, B., Abidi, S.S.R.: Merging Disease-Specific Clinical Guidelines to Handle Comorbidities in a Clinical Decision Support Setting. In: Peek, N., Marín Morales, R., Peleg, M. (eds.) *Artificial Intelligence in Medicine, Lecture Notes in Computer Science*, vol. 7885, pp. 28–32. Springer (2013)
12. Kenning, C., Fisher, L., Bee, P., Bower, P., Coventry, P.: Primary care practitioner and patient understanding of the concepts of multimorbidity and self-management: A qualitative study. *SAGE Open Medicine* 1 (2013)
13. López-Vallverdú, J.A., Riaño, D., Collado, A.: Rule-Based Combination of Comorbid Treatments for Chronic Diseases Applied to Hypertension, Diabetes Mellitus and Heart Failure. In: Lenz, R., Miksch, S., Peleg, M., Reichert, M., Riaño, D., ten Teije, A. (eds.) *Process Support and Knowledge Representation in Health Care, Lecture Notes in Computer Science*, vol. 7738, pp. 30–41. Springer (2013)
14. de Moura, L., Bjørner, N.: Z3: An Efficient SMT Solver. In: *Proceedings of the Theory and Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'08)*. *Lecture Notes in Computer Science*, vol. 4963, pp. 337–340. Springer (2008)
15. Mutasingwa, D.R., Ge, H., Upshur, R.E.: How applicable are clinical practice guidelines to elderly patients with comorbidities? *Canadian Family Physician* 57(7), e253–e262 (2011)
16. Negrini, S., Giovannoni, S., Minozzi, S., Barneschi, G., Bonaiuti, D., Bussotti, A., D'Arienzo, M., Lorenzo, N.D., Mannoni, A., Mattioli, S., Modena, V., Padua, L., Serafini, F., Violante, F.: Diagnostic therapeutic flow-charts for low back pain patients: the Italian clinical guidelines. *Europa Medicophysica* 42(2), 151–70 (2006)
17. Piovesan, L., Molino, G., Terenziani, P.: An ontological knowledge and multiple abstraction level decision support system in healthcare. *Decision Analytics* 1(1), 8 (2014)
18. Salive, M.E.: Multimorbidity in older adults. *Epidemiologic Reviews* 35(1), 75–83 (2013)
19. Shaneyfelt, T., Mayo-Smith, M., Rothwangl, J.: Are guidelines following guidelines? The methodological quality of clinical practice guidelines in the peer-reviewed medical literature. *JAMA* 281(20), 1900–5 (1999)
20. Wilk, S., Michalowski, M., Tan, X., Michalowski, W.: Using First-Order Logic to Represent Clinical Practice Guidelines and to Mitigate Adverse Interactions. In: Miksch, S., Riaño, D., ten Teije, A. (eds.) *Knowledge Representation for Health Care, Lecture Notes in Computer Science*, vol. 8903, pp. 45–61. Springer (2014)